

Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

REMARKS

This application was examined with claims 16- 20 and 21-25. Claims 16-25 remain in the case

The Office Action mailed June 29, 2004 has been carefully considered. Applicants request reconsideration and reexamination of the above-identified application in view of the remarks below, which state Applicants bases for making this request.

The Examiner's presumption concerning inventorship is respectfully noted and is confirmed.

The Examiner's rejection of claims 16-25 under 35 U.S.C. §103(a) as being unpatentable over U.S. Patent No. 6,697,846 [Soltis] in view of U.S. Patent No. 4,761,737 [Duvall] is hereby traversed and reconsideration and reexamination of the claims is respectfully requested in light of the following remarks.

Applicants' independent claim 16 recites a memory storage system having devices organized in physical data blocks for physical storage of data and at least one processor including an operating system having an extent based file system for abstracting file names to the physical data blocks in the devices by assigning an inode to each file. Each inode is adapted to store extents having a field to point to a logical volume. At least two of the extents are direct extents indicating a logical volume containing data blocks. A first direct extent points to first data blocks in the data storage devices and a second direct extent points to second data blocks in the data storage devices. The first direct extent indicates a different logical volume than a second direct extent. Claims 17, 18, 19 and 20 depend from claim 16.

Applicants' independent claim 21 recites a distributed storage system including a global

Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

cache memory with multiple processors and multiple data storage devices coupled to it. Each processor has a local memory for storing an operating system. The devices and processors are capable of communicating by posting messages to each other in the cache memory. Each of the devices includes a processor and local memory storing an operating system. Each operating system has an extent based file system for abstracting file names to physical data blocks in the devices by assigning an inode to each file. Each inode is adapted to store extents having a field which points to a logical volume. At least two of the extents are direct extents indicating a logical volume containing data blocks. A first direct extent points to first data blocks in the plurality of data storage devices and a second direct extent points to second data blocks in the plurality of data storage devices. The first direct extent indicates a different logical volume than a second direct extent. An indirect extent is inserted in the inode between the first and second direct extents, the indirect extent pointing to third physical data blocks in the data storage devices, and at least one extent is written to the third physical data blocks in the data storage devices. The At least one extent written to the third physical data blocks points to fourth physical data blocks. Each of the first, second, and fourth data blocks stores a segment of the file. Claims 22, 23, 24, and 25 depend from claim 21.

Applicants will first discuss the Soltis reference with respect to Applicants' independent claims 16 and the claims which depend from them. The Soltis reference discloses a file storage system suitable for providing users and applications with access to shared data found on storage devices attached directly to shared data found on storage devices attached directly to a network. The file system uses layering techniques to inherit file management functionality from existing systems. Metadata is stored and shared among multiple computers by storing the meta-data as

real-data in regular files of a standard client-server distributed file system. The name space consisting of inode files stored as real-data on the meta-data server acts as the name space for the shared data and file attributes of the inode files are utilized as the file attributes of the shared data. (Soltis, col. 5, line 62 – col. 6, line 5 and Abstract). However, Soltis does not teach or suggest any extent-based method for striping the data from a single file on a single system across multiple logical volumes nor does he show or suggest in any way that two direct extents in a single inode can be used to indicate different logical volumes from one another.

Applicants' respectfully disagree that Soltis teaches a system that each inode adapted to store extents having a field to point to a logical volume. Within any Soltis inode, only one logical volume is directly addressed. Several NAS storage devices form a logical volume to comprise a single, logical device (col. 10, lines 15-22). However, he teaches and shows in his Fig 5 that the data blocks for each file must reside on a single logical volume. Each inode file maintains information pertaining to a single SFS regular file stored on an NAS device (col. 10, lines 53-65). Soltis shows that SFS partitions logical NAS devices into multiple segments in order to exploit parallelism in the network environment (col. 10, lines 35-50). Soltis also shows that all the list extents contained in an inode or set of inodes containing a file indicate physical data blocks in a single logical volume (col. 10, line 53 – col. 11, line 7). It is the physical data blocks themselves, not the extents in the inodes, which may contain pointers to a second logical volume, and even there, all of the pointers contained in the blocks of the first logical volume are all shown pointing to blocks in one single second logical volume, not split among multiple logical volumes. (col. 10, line 53 – col. 11, line 7). In other words, Soltis only teaches that all

the pointers in one inode can point to one logical volume and all the pointers in another inode can point to a second logical volume (Soltis, col. 10, lines 35-50).

Applicants respectfully agree with the Examiner that Soltis does not teach having devices organized in physical data blocks for physical storage of data and at least one processor including an operating system having an extent based file system for abstracting file names to the physical data blocks in the devices and at least two of the extents being direct extents indicating a logical volume containing data blocks, a first direct extent pointing to first data blocks in the data storage devices and a second direct extent pointing to second data blocks in the data storage devices, the first direct extent indicating a different logical volume than a second direct extent. Applicants respectfully point out that, because of this, Soltis cannot and does not teach or suggest any extent-based method for striping the data from a single file on a single system across multiple logical volumes nor does he show or suggest in any way that two direct extents in a single inode can be used to indicate different logical volumes from one another. These are the problems solved by Applicants' invention. Applicants' claim 16 recites that Applicant's inode contains at least two direct extents indicating a logical volume containing data blocks in which segments of a single file are stored, and that the first direct extent indicates a different logical volume than the second direct extent. This feature of Applicants invention is illustrated, for example, in Applicants' Fig 5 and described, for example, at page 7, lines 3-18 of the present disclosure where it is indicated that each direct extent may indicate a logical volume, that the various extents of each inode may map to data blocks of different logical volumes, and these extents may map different segments of a single abstract file to different ones of the drives and to different physical disks and partitions thereof.

Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

The deficiencies of Soltis with respect to Applicants' invention are not overcome by Duvall. The Duvall reference discloses a UNIX file system managed in a virtual machine environment. Duvall's system manages the allocation of virtual address space in the system. Applicants' acknowledge that Duvall does teach data block addresses and plurality of sequence of blocks, (col. 2, lines 53-68) and Applicant's also acknowledge that Duvall discloses an inode structure (col. 2, lines 62 -67 and col. 3, lines 1-6). However, Duvall does not teach or disclose the logical volume concept. Indeed, the invention in Duvall involves the use of a map page range service for mapping files directly to a physical address (col 7, lines 10-68, col 8, lines , col 19, lines 60-68, col 20, lines 1-40). Thus, Duvall does not teach, disclose or suggest any provision for writing different direct extents to addresses in different logical volumes, as discussed in Applicant's description and specifically recited in Applicant's claims 16 and 21. Nor does Duvall teach, disclose or suggest any reason for such assignment. Thus, Duvall cannot teach, disclose or suggest any inode in which a first direct extent specifies a first logical volume and a second direct extent specifies a second logical volume.

Applicants respectfully submit that Even if Soltis and Duvall were combined, neither of the two, alone nor in combination, teaches or suggests Applicants' invention as described in claim 16. There is no suggestion in either reference of how extents within a single inode could map to multiple logical volumes. It remained to Applicants to invent a solution to this problem. Moreover, since Soltis teaches file management in a distributed, shared file system in a client-server environment and Duval teaches memory management in system using virtual memory, there would be no reason to combine the two references, or to apply either of them to data management. Applicants respectfully submit that claim 16 is allowable.

Since claims 17, 18, 19 and 20 depend from claim 16, Applicants respectfully submit that these claims are allowable for at least the same reasons as for claims 16.

Applicants will now discuss the Soltis reference with respect to Applicants' independent claim 21 and the claims which depend from it. Applicants respectfully agree with Examiner's statement that Soltis does not show a distributed storage system including a plurality of processors and a plurality of devices coupled to a global cache memory, each of the processors having a local memory storing an operating system. In Soltis, it is a first host computer with a first operating system which creates the first inode and data blocks in mass storage memory, and the second host computer with a second operating system which creates the second inode and second data blocks. While both of Soltis' inodes are stored in the storage memory, neither is created by any part of the mass storage system. The file systems of the hosts create the inodes using metadata suitable to that particular host. (Soltis, col. 8, lines 27-52 and col. 10, line 51-col. 11, line 43). By contrast, in accordance with applicants' invention, it is the extent based file system resident in the operating system of each component (processor and device) of the distributed storage system itself which creates the inode and the extents in the inode in order to pass messages within the distributed storage system without dealing with an intermediate abstraction layer for handling logical volumes. Structural features of the distributed storage system of Applicants' invention are described, for example in Fig 3 and at page 5, lines 10-33 of the present disclosure where it is indicated that each processor and driver has a permanent memory storing microcode which includes an operating system with a file system and the cache memory provides a symmetric environment for communications between the processors and drivers, which send requests to and respond to requests from others by writing messages in

Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

predetermined locations of the global cache memory. I-node features of Applicants' invention are described, for example in Figs 10 and at page 11, lines 23-30 and page 13, lines 1-10 of the present disclosure where it is indicated that the operating system of the device or processor component of the distributed storage system assigns an inode to the file by writing the inode address and the file name to a row in the directory. All of these features are specifically recited in Applicants' claim 21.

The deficiencies of Soltis with respect to claim 21 are not overcome by the addition of Duvall. Applicants respectfully disagree with Examiner's statement that Duvall shows a mass storage system with a plurality of processors coupled to the global cache memory, each processor having a local memory for storing an operating system and the processors capable of communicating by posting messages to each other in the cache memory, each of the devices including a processor and local memory storing an operating system, each operating system having an extent based file system for abstracting file names to physical data blocks in the devices by assigning an I-node to each file, each I-node adapted to store extents having a field to point to a logical volume. In Duvall, the file system is part of an operating system executing on a host. The host is in communication with a secondary storage system, but the file system, not the secondary storage system itself, creates the inode. Duvall's system provides a method for use by a page segmented virtual memory system to increase the number of virtual pages that were initially assigned in the segment and the number of corresponding storage blocks on a secondary storage device (claim 1 and col. 8, lines 22-36). System organization features of Applicants' invention are described, for example in Fig 3 and at page 5, lines 20-26 of the present disclosure where it is indicated that the cache memory provides a symmetric environment for

communications between the processors and drivers, which send requests to and respond to requests from others by writing messages in predetermined locations of the global cache memory. Extent features of Applicants' invention are described, for example, in Figs 6A and 6B and at page 7, lines 12-24 of the present disclosure where it is indicated that each extent includes an address pointer field which indicates both a logical volume and a physical offset of a data block in the logical volume.

Applicants again respectfully agree with the Examiner that Soltis does not teach having devices organized in physical data blocks for physical storage of data and at least one processor including an operating system having an extent based file system for abstracting file names to the physical data blocks in the devices and at least two of the extents being direct extents indicating a logical volume containing data blocks, a first direct extent pointing to first data blocks in the data storage devices and a second direct extent pointing to second data blocks in the data storage devices, the first direct extent indicating a different logical volume than a second direct extent. As Applicants have pointed out in connection with claim 16, neither Soltis nor Duvall nor any combination of the two teach or suggest any extent-based method for striping the data from a single file on a single system across multiple logical volumes nor does he show or suggest in any way that two direct extents in a single inode can be used to indicate different logical volumes from one another. These are the problems solved by Applicants' invention. Applicants' claim 21 recites that Applicant's inode contains at least two direct extents indicating a logical volume containing data blocks in which segments of a single file are stored, and that the first direct extent indicates a different logical volume than the second direct extent. This feature of Applicants invention is illustrated, for example, in Applicants' Fig 5 and described, for example,

Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

at page 7, lines 3-18 of the present disclosure where it is indicated that each direct extent may indicate a logical volume, that the various extents of each inode may map to data blocks of different logical volumes, and these extents may map different segments of a single abstract file to different ones of the drivers and to different physical disks and partitions thereof.

Applicants also respectfully agree with the Examiners statement that Soltis does not disclose writing a plurality of extents to an inode assigned to a file, writing to first and second data blocks, inserting an indirect extent between two direct extents and writing at least one extent to the data block indicated by the indirect extent, as claimed by Applicants. Applicants' claim 21 recites inserting an indirect extent in the inode between the first and second direct extents. These features of Applicants' invention are described, for example, in Figs 8A and 8B and page 8, lines 22-33 and page 9, lines 1-7 and specifically recited in Applicants' claim 21. It is not possible to insert an indirect extent between two direct extents in the system disclosed by Soltis because in Soltis' system indirect extents and direct extents must occupy different positions in the inode array (Soltis, col. 16, lines 27-45). Soltis neither teaches, discloses or suggests an insertion as claimed by Applicants in their claim 21. The deficiencies of Soltis with respect to Applicants' invention are not overcome by Duvall. With regard to claim 21, Applicants respectfully disagree with the Examiners statement that Duvall teaches writing a plurality of extents to an inode, writing data to first and second data block, inserting an indirect extent in the inode between the two direct extents, as claimed by Applicants, and writing at least one extent to the third physical data block, the extent pointing to a physical data block storing a segment of the file. As Applicants' have pointed out in connection with Soltis, claim 21 recites inserting an indirect extent in the inode between the first and second direct extents. It is not possible to insert an

Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

indirect extent between two direct extents in the system disclosed by Duvall because in Duvall's system indirect extents and direct extents must occupy different positions in the inode array (Duvall, col 2, lines 62-68 and col 3, lines 1-6). In Duvall's system, each group of contiguously allocated blocks is summed, and the count recorded in the field adjacent to the starting block entry in the Map Page Range structure and discontinuous blocks are reflected in discrete entries in the Map Page Range structure (Duvall, Fig 7c and col. 19, lines 42-59). Applicants respectfully submit that this claim is allowable.

Furthermore, with respect to dependent claims 22-25, Applicants respectfully submit that since these claims depend from claim 21, they are allowable for at least as the same reasons as for claim 21.

The remaining references cited by the Examiner have been reviewed with respect to the claims remaining in the case, and are not considered to adversely affect patentability of these claims.

Accordingly, it is respectfully submitted that Applicants' have overcome the Examiner's rejections with respect to all remaining claims and put them in condition for allowance. Applicant requests reconsideration and reexamination of the above-identified application in view of the remarks above, which state Applicant's bases for making this request.

In the event the Examiner deems personal contact desirable in the disposition of this case, the Examiner is invited to call the undersigned attorney at (508) 293-7998.


Applicant: Preston F. Crow, *et al.*
U.S.S.N.: 10/720,629
Filing Date: November 24, 2003
EMC Docket No.: EMC-99-027DIV1

Please charge all fees occasioned by this submission to Deposit Account No. 05-0889.

Respectfully submitted,

Dated:

Oct 29, 2004



Penelope S. Wilson, Esq. (Reg. No. 29,751)
Attorney for Applicants
EMC Corporation
Office of General Counsel
176 South Street
Hopkinton, MA 01748
Telephone: (508) 293-7998
Facsimile: (508) 293-7189